A Forestry Expert Package - The Lake Traverse Study

F. Gougeon
Petawawa National Forestry Institute • Information Report PI-X-108

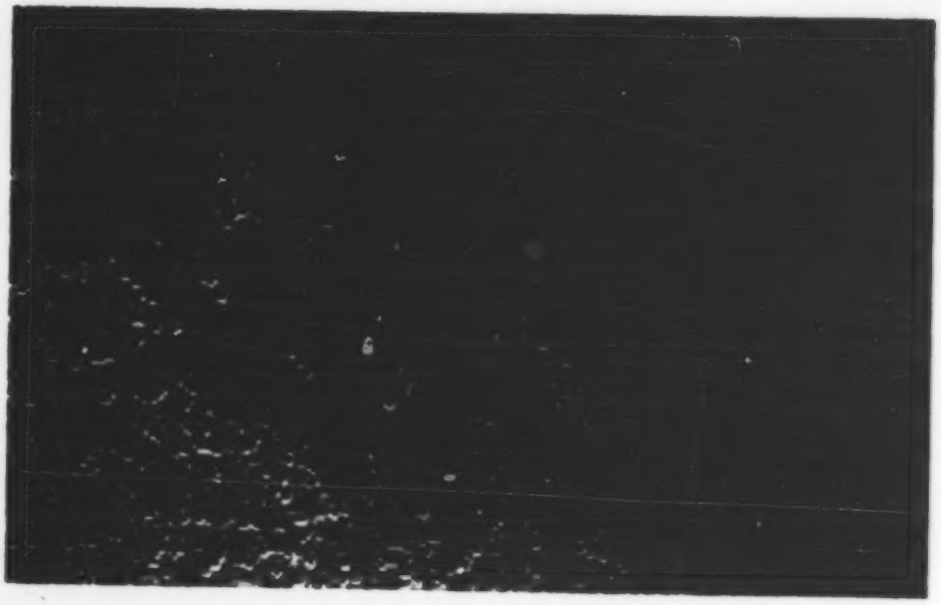

Forestry Canada Forêts Canada

Canada

92-04096/

PETAWAWA NATIONAL FORESTRY INSTITUTE

In common with the rest of Forestry Canada, the Petawawa National Forestry Institute has as its objective the promotion of better management and wiser use of Canada's forest resource to the economic and social benefit of all Canadians. Objectives of program activities carried out at the Institute support this goal through discovery, development, demonstration, implementation, and transfer of innovations. Because it is a national institute, particular emphasis is placed on problems that transcend regional boundaries or that require special expertise and equipment that cannot be duplicated in Forestry Canada regional establishments. Such research is often performed in close cooperation with staff of the regional centres, provincial forest services, and the forest industry.

Research initiatives and technical services at the Institute encompass five major activities:

FOREST GENETICS AND BIOTECHNOLOGY — Integrates projects in tree genetics, soil microbiology, micropropagation, molecular genetics, meteorology, and seed research. It also includes the client services and seed bank operations of the National Tree Seed Centre, a long-standing program with extensive international affiliations.

FOREST MANAGEMENT SYSTEMS — This program integrates projects in fire, remote sensing, modelling, growth and yield, and forest pest management to provide research and development for the formulation and demonstration of forest management systems.

NATIONAL FOREST RESOURCE STATISTICS — Provides biological, technical, and socioeconomic information on Canada's forest-based resources. The program involves progressive development of databases and establishment of new databases and software in support of policy development in forestry. The Forest Inventory Program collates information on the forest resource at a national level, maintains the Canadian Forest Resources Data System, and prepares the national forest inventory.

COMMUNICATIONS — Integrates activities of the library, public awareness, information, and editing and publications projects. The Institute is visited by more than 20 000 people every year. There is a Visitor Centre for the public, self-guided tours, and an extensive education project. The national repository of all scientific and technical publications of the Forestry Canada and the principal Forestry Canada publications distribution centre are both located at PNFI.

THE RESEARCH FOREST — Besides natural stands manipulated in a variety of ways for silvicultural research, the 100 km^2 Petawawa Forest contains extensive areas of plantations dating back six decades. Research plantations are a source of growth and yield data derived from cultural experiments, and they are becoming valuable for pedigreed genetic materials for micropropagation and molecular genetics studies. The forest also offers opportunities for short- and long-term testing of forest management strategies.

A Forestry Expert Package - The Lake Traverse Study

F. Gougeon

Information Report PI-X-108
Petawawa National Forestry Institute
Forestry Canada
1991

©Minister of Supply and Services Canada 1991
Catalogue No. Fo46-11/108-1992E
ISBN 0-662-19317-2
ISSN 0706-1854
Printed in Canada

Copies of this publication may be obtained free of charge from the following address:

Forestry Canada
Publications Distribution Centre
Petawawa National Forestry Institute
Chalk River, Ontario
K0J 1J0

Telephone: 613-589-2880

A microfiche edition of this publication may be purchased from:

Micromedia Ltd.
Place du Portage
165, Hôtel-de-Ville
Hull, Québec
J8X 3X2

Cette publication est également disponible en français sous le titre
Un Progiciel Expert en Foresterie : étude du lac Traverse.

CANADIAN CATALOGUING IN PUBLICATION DATA

Gougeon, François, A.

A forestry expert package : the Lake Traverse study

(Information report ; PI-X-108)
Includes an abstract in French.
Includes bibliographical references.
ISBN 0-662-19317-2
DSS cat. no. Fo46-11/108-1992E

1. Forest surveys -- Ontario -- Computer programs.
2. Forest management -- Ontario -- Computer programs. 3. Expert systems (Computer science).
I. Petawawa National Forestry Institute.
II. Title. III. Title: The Lake Traverse study.
IV. Series: Information report (Petawawa National Forestry Institute) ; PI-X-108.

SD387.S86G68 1992 634.9'2'09713 C92-099545-4

Contents

v Abstract/Résumé

1 Introduction

2 The Forestry Expert Package

6 Study Area Description and Available Data

6 Landsat Image Classifications

8 Testing the FEP on the Lake Traverse Area

11 FEP weaknesses

14 Conclusion

15 Acknowledgments

15 References

Tables

2 1. A typical class transition table (counts of pixels in transitions from a specific summer 1974 (S74) class to a specific summer 1976 (S76) class, classes being numbered 1 to 8)
6 2. 1981 Region descriptions for Figure 2
8 3. Decision and summer class descriptions
8 4. Winter class descriptions
8 5. Classification accuracies
11 6. Percentage of pixels in the same class when comparing two decision images
13 7. Decision image accuracies
13 8. Percentage of pixels in the same class when comparing two classifications

Figures

4 1a. The Forestry Expert Package flowchart
5 1b. The Forestry Expert Package flowchart (cont.)
7 2. Reforestation activities south of Lake Traverse, Algonquin Park, Ontario (see Table 2 for a description of regions A-K)
12 3. A Landsat image unsupervised classification of the summer of 1974 and/or Decision Image One produced by the FEP (subset of full image showing only areas where activities took place)
12 4. Decision Image Two which indirectly reflects the combination of the summer 1974 and winter 1976 classifications
12 5. Decision Image Three produced by the addition of the summer 1976 classification
12 6. Decision Image Four produced by the addition of the winter 1981 classification
12 7. Decision Image Five - produced by the addition of the summer 1981 classification

ABSTRACT

The detection of ongoing changes in forested areas monitored with satellite images is traditionally accomplished with multi-date band ratioings, enhancements, and classifications. The Forestry Expert Package (FEP) uses a completely different method and is an example of a new philosophy in image analysis: an artificial intelligence approach. It uses single-date Landsat classifications as input data, compares them, and decides on the significance and cause of detected changes based on rules derived from forestry knowledge.

For the present experiment, classifications of geometrically corrected images of different years and different seasons were obtained with an unsupervised classifier for a test area near Lake Traverse in Algonquin Park, Ontario. Several tests were run to determine if the FEP would detect logging activities and the subsequent regeneration known to have occurred in this area. A secondary objective was to see if any improvement in the classification of forest types of the non-changing areas could be achieved over conventional single data classifications.

The FEP was found to be good at detecting changes in the forest due to logging, and moderately good at monitoring the regrowth of conifers. However, the FEP is not yet at a stage where it could be of use as a forest management decision-making tool.

RÉSUMÉ

La détection des changements en régions boisées observées à l'aide d'images satellitaires est traditionnellement accomplie par des rapports de bandes, des améliorations d'images et des classifications multidates. Le Progiciel Expert en Foresterie (PEF) utilise une méthode complètement différente et est un exemple d'une nouvelle philosophie de l'analyse des images: une approche basée sur l'intelligence artificielle. Il utilise des classifications unidates d'images Landsat comme données d'entré, les compare, et décide de l'importance et de la cause des changements perçus en se fondant sur des règles dérivées du savoir forestier.

Pour cette expérience, des classifications d'images corrigées géométriquement d'une région près du lac Traverse au Parc Algonquin, en Ontario, ont été obtenues par un processus de classification non-supervisé pour des années et des saisons différentes. Plusieurs tests ont été faits pour déterminer si le PEF détecterait les coupes et la régénération subséquente qui ont marqués la région. Un objectif secondaire fut de voir si une amélioration de la classification des types forestiers des régions inchangées pourrait être obtenue par rapport aux classifications unidates conventionnelles.

Le PEF fut estimé bon pour détecter les changements du couvert forestier dûs à la coupe, et modérément bon pour suivre la régénération des résineux. Toutefois, le PEF ne s'est pas encore rendu au point d'être utile en tant qu'outil de prise de décisions en gestion forestière.

INTRODUCTION

Forest managers and planners are increasingly relying on computerized resource inventories (Gray and Nietmann 1989, Power 1987), change monitoring (Gillis 1988), and modelling and decision support tools (Newnham 1987, Moore and Lockwood 1990). These tools need to be "flexible" so that decision-makers can bring their forestry experience and specific knowledge of the area under scrutiny to bear on the problems at hand. The desired system should also be "user friendly", requiring only a minimum amount of digital image analysis (DIAS) or geographic information system (GIS) expertise on the part of the user. Artificial intelligence-based systems have a good potential for providing these features, thus the growing movement towards the use of expert systems in forestry (Kourtz 1990).

During the early 1980s, some thought was given to applying artificial intelligence to remote sensing image analysis (NASA 1979, Tinney et al. 1983, Mooneyham 1983). Researchers at Kyoto University (Nagao and Matsuyama 1980) went even further and created a sophisticated system that first segmented aerial photographs and then identified the various segments based on some keep basic properties (average grey levels, texture, shape). Segments could even be grouped together to form objects (e.g., a house). Other applications have been developed subsequently, such as the Canada Centre for Remote Sensing's SHERI system (Goodenough et al. 1987, 1990a, 1990b), which includes numerous specialized satellite image analysis expert systems under one umbrella, and Carnegie-Mellon University's MAPS system (McKeown 1984, McKeown et al. 1990), which deals mostly with aerial images of urban areas.

Traditionally, in digital remote sensing and computerized image analysis, efforts toward the detection of changes over time have followed the routes of multi-date image enhancement, classification, and band ratioing (Hall et al. 1984, Banner and Lynham 1981). Little work has been done on post-classification change detection (Howarth and Wickware 1981). However, as forest resource management becomes more computerized and the use of satellite imagery to update databases more popular (Hegyi and Quenet 1983, Pilon and Wiart 1990), the preferred medium for change detection and monitoring could become databases containing thematic information, rather than raw satellite (e.g., Landsat, SPOT, MOS) data.

Digital remote sensing is generally not yet well integrated in the production of forest inventories. These inventories are still mostly produced in conventional ways, with the resulting information painstakingly transferred to modern GISs. The integration of DIASs and GISs will involve the "full" integration of raster- and vector-based information (i.e., not only for visualization purposes). Until sophisticated GISs or DIASs exist that can truly and invisibly deal with information from both of these domains, one has to decide which domain is to be used to perform comparisons for change detection and updating purposes. Of course, GISs have tools to compare information in their native vector form and, sometimes, tools to convert raster information to vectors. However, vectors are inherently more complicated for comparison operations. The Forestry Expert Package (FEP) is a prototype of a system that uses artificial intelligence to compare thematic information in raster form. As a first step in that direction, it compares classified Landsat images taken at different dates to mainly detect and categorize changes that may have occurred.

The concepts leading to the creation of the FEP can be traced to 1977 (Goldberg and Kourtz 1977). In 1982, the FEP was developed for the Petawawa National Forestry Institute's Digital Remote Sensing Project by S.C.S. Consultants Inc. via an unsolicited proposal (UPS-392). Several articles (Goldberg et al. 1983a,b) have already reported on some of its features and capabilities. The present report differs from its predecessors in two ways: 1) It describes the FEP at a more conceptual level and outlines its uses and benefits in more general terms, addressing foresters rather than the artificial intelligence community; 2) Some of the weaker features of the current version of the FEP are compensated for, to show the validity of the philosophical approach and the potential benefits of such systems.

The present report describes how the Forestry Expert Package can be used to detect changes in forested regions monitored via Landsat image classifications. It reviews the results of testing the FEP on an intensively-managed forested area of Algonquin Park, Ontario. A discussion of possible improvements concludes the report.

The Forestry Expert Package

The FEP is an expert system that assists the monitoring of a region through time via digital remote sensing data. It facilitates the detection of changes and, hopefully, ensures more reliable information about the areas that have not changed. It uses a series of chronologically ordered satellite image classifications of the region (thematic images in which forest stands are colour-coded according to their content) as input data. These classifications, of geometrically corrected scenes of different years and different seasons (winter and summer), are compared with each other to detect changes that have occurred over time. For each sequential introduction of a new classification, the FEP produces a *decision image*[1], which is essentially the most up-to-date categorization (i.e. classification) of the region. Both the input classifications and the subsequent decision images are associated with *reliability images*, conveying the reliability of a class assignment or decision for each pixel (or picture element, corresponding here to an area of 50 m x 50 m on the ground) in the image. Figure 1 is a block diagram depicting the FEP's operations.

When a new classification is introduced, it is compared with the previous classification of the same season in order to generate a class-based summary table of differences on a pixel by pixel basis. This means that each pixel of one classification is compared with its corresponding pixel (i.e. same position) in the other classification and that any change in class assignment between the two pixels is recorded. Classes and class sequence are assumed the same for both classifications. At the end of the comparison, a *class transition table* (e.g., Table 1) is produced summarizing the number of transitions from each class in the first classification to each class in the other. The diagonal elements of the table are counts of the number of pixels that have not changed class assignment. Similarly, a *reliability transition table* is also produced indicating for each class transition the percentage of pixels that are more reliably classified in the second classification than in the first. Here, classification reliability quantifies how typical a pixel is within its class.

In the class transition table, some of the transitions reflect *real changes* that occurred in the region of interest during the time period under consideration. Unfortunately, most do not. Indeed, in any direct pixel-by-pixel comparison of Landsat classifications, there is always a substantial amount of *artificial changes* recorded. These are created by factors such as georeferencing errors, classification imprecision, the presence of haze, and the relative shifting of class assignments. In order to eliminate these changes from consideration by the decision-making part of the FEP or by the user, a threshold-based *filtering mechanism* is used. Only the changes corresponding to transitions that are judged significant by the filtering mechanism will have any possibility of modifying current decisions.

[1] Expressions that are important to a clear understanding of the FEP will be in *italics* when first introduced and defined.

Table 1. A typical class transition table (counts of pixels in transitions from a specific summer 1974 (S74) class to a specific summer 1976 (S76) class, classes being numbered 1 to 8)

	Summer 1976 (S76)								S74
S74	1	2	3	4	5	6	7	8	Total
1	3394	0	1	2	16	78	536	0	4027
2	0	3986	5638	121	250	9	0	0	10004
3	0	272	9201	1520	2952	54	27	0	14026
4	7	15	2943	2436	6355	438	96	0	12290
5	10	13	847	4319	9641	2159	702	0	17691
6	24	5	180	2066	5060	8731	4642	0	20708
7	330	0	71	132	888	5236	9136	0	15853
8	4	62	1280	34	4043	806	73	0	6302
S76 Total	3769	4353	20161	10630	29205	17511	15212	0	

The filtering mechanism involves, first, a comparison of the transition table with a user-entered table of expected transitions, the *class transition threshold table*. Any class transition that occurs more frequently than expected is judged significant. Secondly, even transitions that do not occur frequently enough can be judged significant if their overall increase in reliability is sufficiently high. Thus class transitions for which the value accumulated in the reliability transition table are above the user-entered threshold stored in the *reliability transition threshold table* are also judged significant. In other words, this filtering mechanism attempts to judge the significance of the class transitions first on a quantity basis, then on a quality basis.

Once the significant class transitions are known, the FEP recompares the new classification with the old one of the same season on a pixel by pixel basis (see Figure 1b). First, it checks for class differences between the two corresponding pixels (the detection section). If a difference is not detected, it is assumed that nothing has changed in the forest during the time interval between the two classifications. The previous FEP decision is maintained and its reliability increased. This increase is justified by the fact that new information (i.e., the new classification) corroborates the previous information on which the previous decision was based. If a class difference between the two corresponding pixels is detected, the filtering mechanism is invoked.

The filtering section checks whether the encountered difference in a pixel's class corresponds to a class transition previously judged significant (based on the criteria outlined above). If that is not the case, the old decision is maintained, but its reliability is lowered to reflect the possibility that their might be some conflicting information about this forest location. Only differences in classification that correspond to significant transitions are passed on to the decision-making mechanism. A significant transition (because of its frequency or its reliability) usually implies that a real change in ground-cover type occurred during the time period between the two classifications. Consequently, a decision should be made as to the nature of this change in the forest. This is the function of the decision-making section of the FEP.

A new decision must be assigned for each pixel where a significant transition has occurred. The decision is made based on the classes assigned to this pixel in the currently introduced classification, the previous classification of the same season, and the previous decision image (the past assessment of the situation). It is made based on previously user-entered rules, such as:

> "If a pixel was considered to represent a soft-wood area in the previous classification and in the previous decision image, and it is categorized 'open-area' in this classification, then the area was probably logged, so make a decision that it is now a 'recent cut' area."

This part of the package can be conceptualized as a rule-based expert system (Hayes-Roth et al. 1983) where all the rules have exactly the same syntax. Each rule consists of the conjunction of three conditions (old class, old decision, and new class) in the premise, producing a single action, namely, the selection of a specific decision. Because these rules must be applied to a sizable number of pixels, they are actually stored and used as a multidimensional lookup table.

The changes for which decision rules are not yet defined are automatically presented to the user for identification. The user is expected to enter a new rule after careful analysis of the situation. To prepare this new rule, the user can look at the *anomaly image* (a temporary image), which shows all the pixels undergoing a significant transition for which no rule exists. Knowing the spatial position of the troublesome pixels (their location within the region being monitored) helps the user better understand the situation and create the appropriate rule. The new rule will be recorded and used automatically by the system each time an identical situation is encountered. Over time, by learning from the user, the system becomes "knowledgeable" enough to make all of the decisions.

In summary, the decision rules are only invoked for a pixel exhibiting a class difference (transition) between a new and an old classification of a given season and for which the class transition was judged significant. Otherwise (no difference or non-significant transition), the old decision is kept and becomes the new decision for that pixel position in the new decision image. Whether the decision rules are used or not, the FEP will assign a reliability factor to all of its decisions. For a pixel without a classification

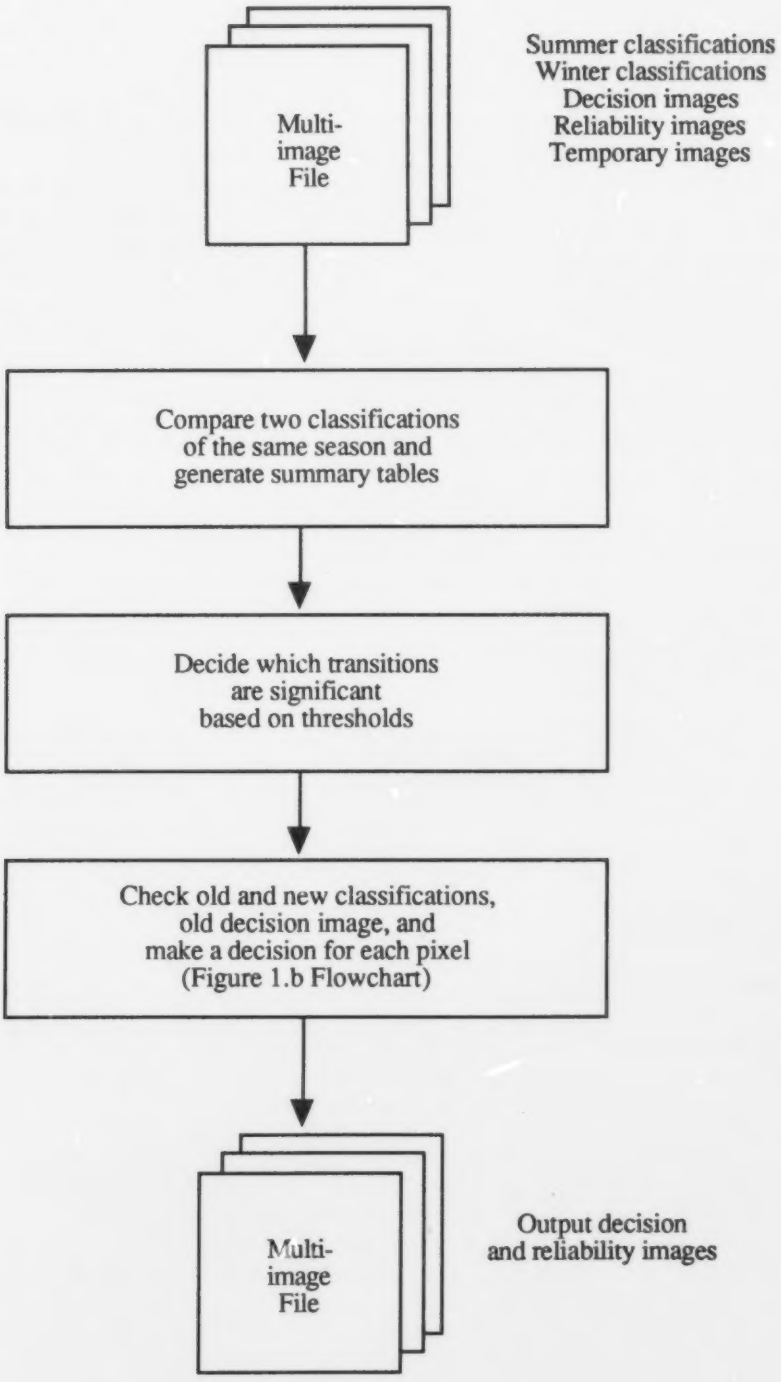

Figure 1a. The Forestry Expert Package flowchart

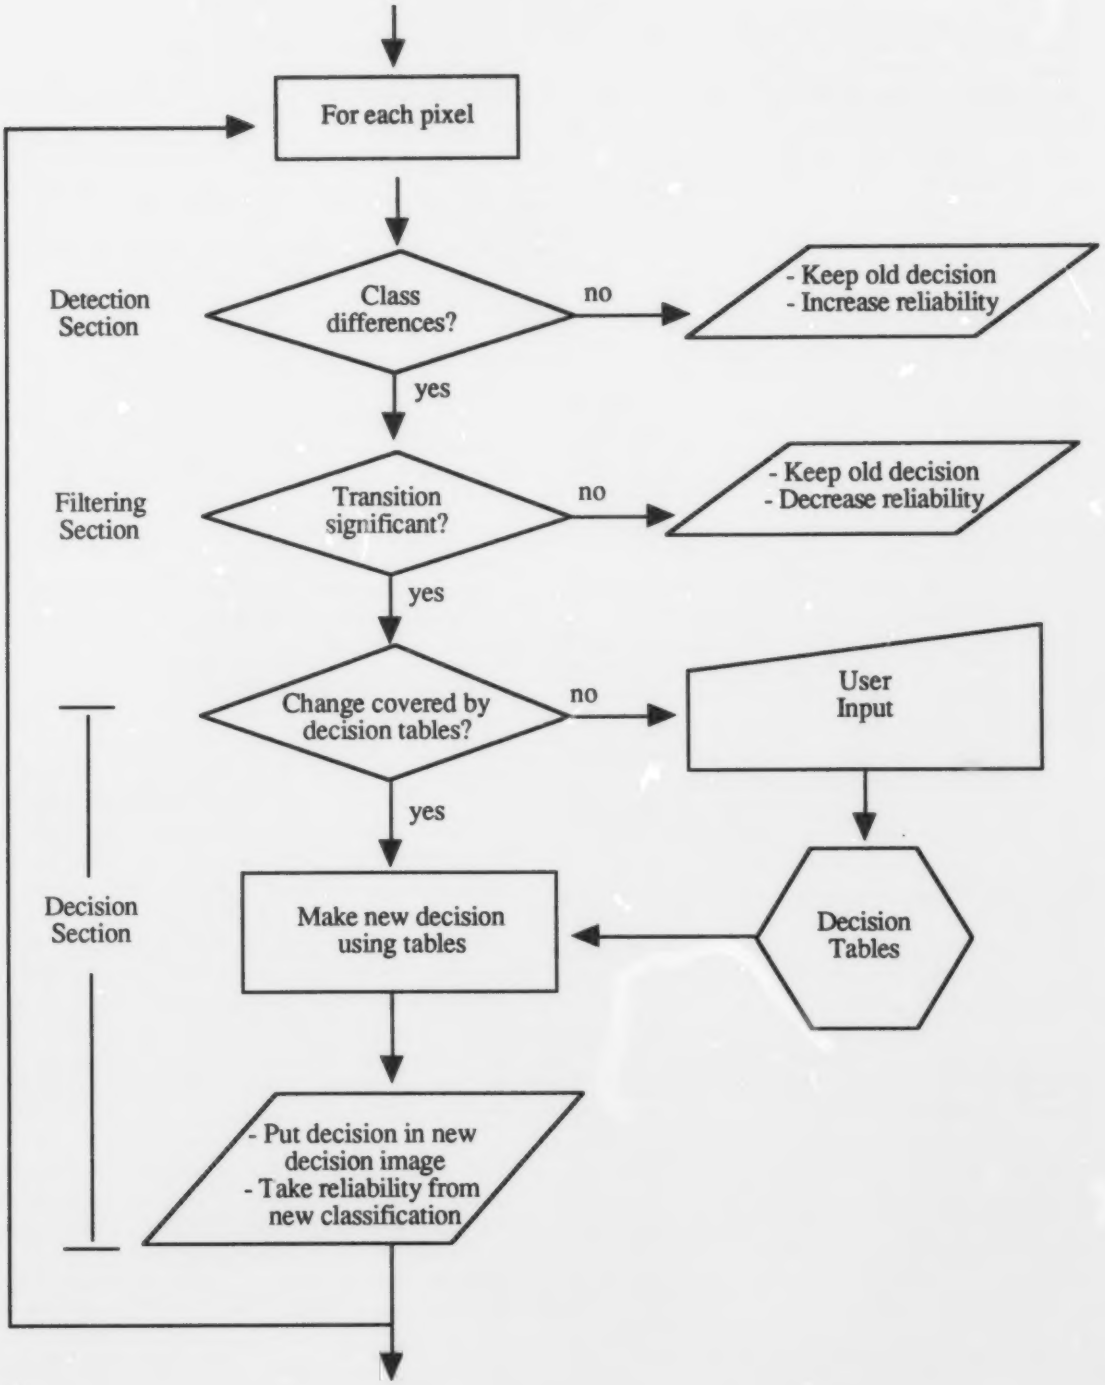

Figure 1b. The Forestry Expert Package flowchart (cont.)

difference, the new decision is the same as the previous decision but is now judged more reliable. That is, its reliability is arbitrarily increased. The previous decision is also kept for a pixel that has been affected by a non-significant transition, but its reliability is lowered. Finally, when a decision has been taken via the decision rules whether newly entered or not, it becomes part of the new decision image and its reliability is assigned to that of the newly introduced classification pixel because this was the driving force behind making the new decision. In this way, a reliability image is produced from each decision image, and both are produced, pixel by pixel, each time a new classification is introduced to the system.

Study Area Description and Available Data

The Forestry Expert Package was tested on a forested area centered on Lake Traverse in Algonquin Park, Ontario. This area is composed mostly of 60- to 100-year-old stands, with varying proportions of hardwoods and softwoods. The area was monitored using geometrically corrected summer and winter Landsat MSS images from 1974 to 1981. The 1978 Ontario Ministry of Natural Resources forest resource inventory (FRI) done with 1977 aerial photographs was available, as well as black & white and colour aerial photographs for various dates and some field samples. Logging activities and subsequent replanting efforts were of great interest and were well documented (see Figure 2 and Table 2 for a summary). These changes in forest cover corresponded to about six percent of the area under observation.

Table 2. 1981 Region descriptions for Figure 2

Region	Age of regrowth (yrs)	Predominant height (m)	Stocking (stems/ha)
A	19	6	2000
B	8	1.5	6000
C	6	1	5000
D	7	1	2800
E	5	0.8	2800
F	4	1	3200
G	6	1.2	2800
H	9	2	14000
I	9	2	17000
K	24	10	N/A

Landsat Image Classifications

It was decided that the classifications of the various Landsat images should be done in such a way as to be as automatic and unbiased as possible. The intent was to ensure that the process would be the same for all images, that classification accuracies would be of the same order, and that the situation would simulate a production environment.

The accuracy of a quick and automatic unsupervised classification is generally of the order of 50 to 70 per cent. It can usually be improved to the 70 to 80 per cent range by further refinements and the use of *a priori* knowledge about the scene being processed. Due to diminishing returns, these efforts can be very time consuming. In a production run of the FEP, the user may not have the time nor the *a priori* knowledge to improve the classification. In addition, an anticipated advantage of the FEP is that it will improve the accuracy of the non-changing areas in the final decision image beyond the best accuracy possible with a single date classification. In other words, the quality lost to speed and convenience at the single classification level is expected to be partially compensated for by the quantity of classifications done of the same area.

A multispectral classifier (Goldberg and Shlien 1976, Goldberg et al. 1982b) was used to obtain unsupervised classifications for summer and winter Landsat MSS images of 1974, 1976, and 1981. This classifier is based on isolating the most important peaks in a four dimensional histogram of the multispectral data (Landsat MSS gathers images for each of four visible to near-infrared spectral bands). It uses various automatically generated thresholds in order to separate the most distinctive classes. As mentioned above, the process was carried in such a way as to produce classifications as automatically as possible. That is, the classifications were done solely based on statistical separations and regroupings to obtain the desired number of classes, without the benefit of displaying any intermediate classification. No effort was made to improve a classification once obtained. Due to significant modifications of the spectral values by prevalent haze, a good classification of the 1974 winter image could not be achieved and that season had to be dropped from the FEP test.

The classifications obtained for the Lake Traverse area (300 x 340 pixels, or 15 km x 17 km) consisted of eight classes for the summer scenes and four classes for the winter scenes (see Tables 3

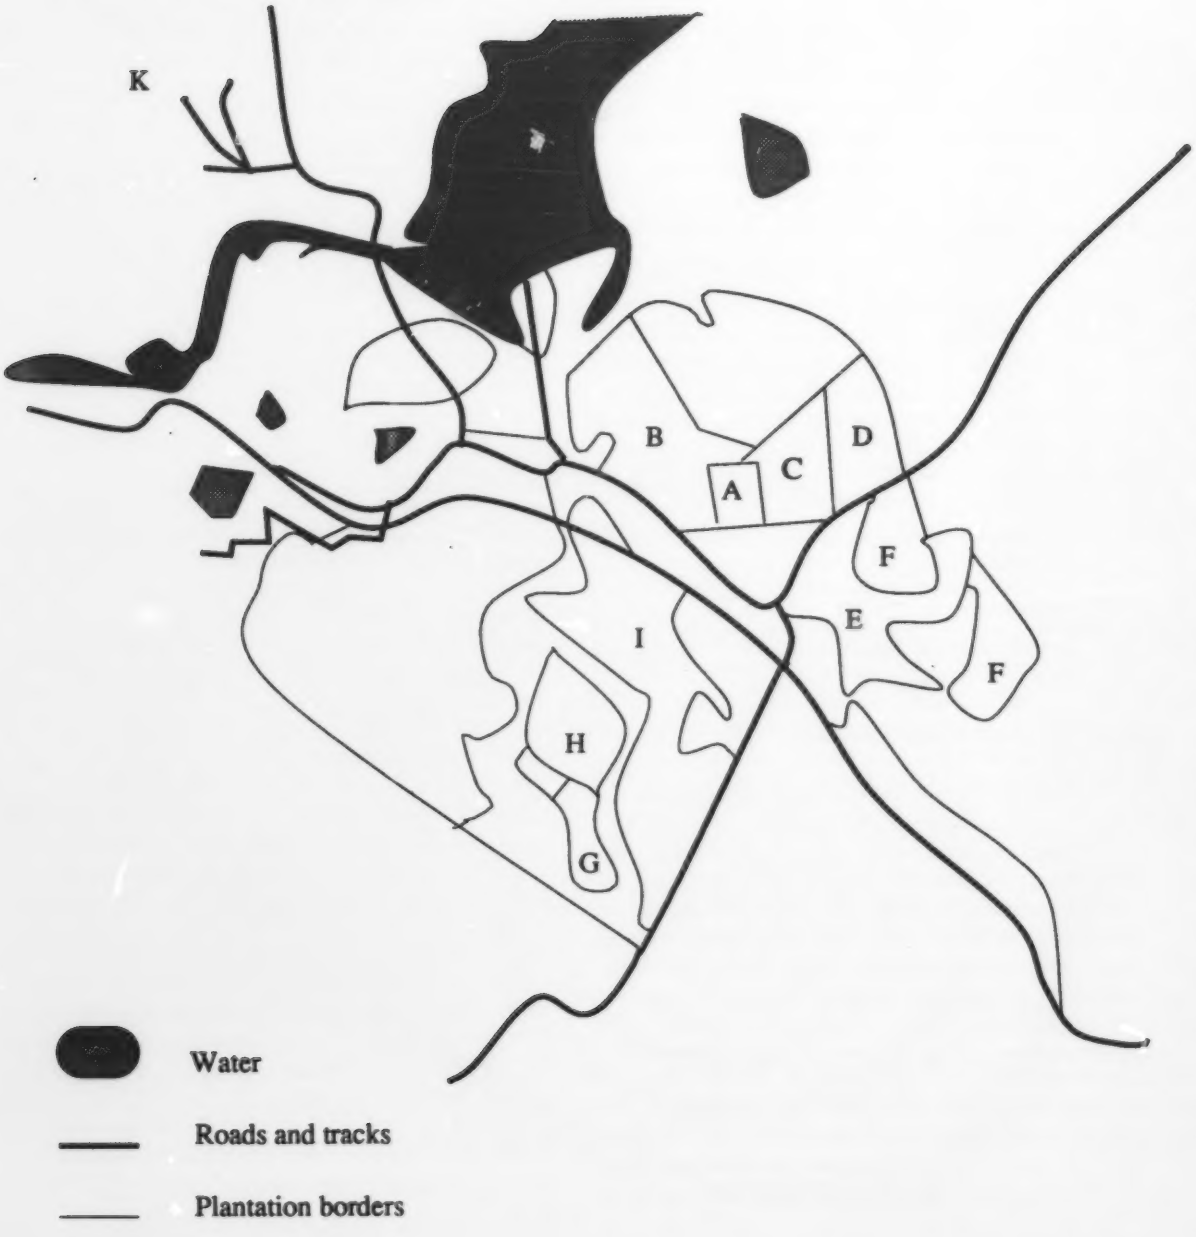

Figure 2. Reforestation activities south of Lake Traverse, Algonquin Park, Ontario (See Table 2 for a description of regions A-K)

and 4). The winter classes were snow, hardwood, mixedwood, and softwood. The summer classes were more numerous and followed a rather precise breakdown not usually attempted with Landsat MSS data: water, above 80% hardwood, 70-80% hardwood, 60-70% hardwood, mixedwood/open areas, 60-70% softwood, above 70% softwood, and recently cut areas. It was felt that these narrowly-defined classes would be needed to verify if the FEP showed an improvement in classification accuracy over conventional classifications.

Table 3. Decision and summer class descriptions

Decision class	Summer class	Colour	Description
1	1	mBlue	Water
2	2	mGreen	Above 80% hardwood
3	3	dGreen	70-80% hardwood
4	4	dBlue	60-70% hardwood
5	5	Yellow	Mixedwood and/or open
6	6	Orange	60-70% softwood
7	7	Brown	Above 70% softwood
8	8	Red	Cut and/or open
9		Mauve	Old cut
10		lBlue	General regeneration
11		dBrown	Softwood regeneration
12		White	Pure open
13		llBlue	Man-used

Table 4. Winter class descriptions

Winter class	Colour	Description
1	Blue	Snow covered areas
2	Green	Hardwood areas
3	Yellow	Mixedwood areas
4	Brown	Softwood areas

Given the rather precise class breakdown used in the summer classifications and the fact that misclassification to a spectrally neighbouring class is not very damaging to the decision-making process in the FEP, it is possible to view classification accuracies from two points of view: strict and general. Here, *strict accuracy* refers to the accuracy figures obtained from the diagonal elements of a contingency table (or confusion matrix), comparing the classes assigned by the classifier with those known to be true on the ground for pixels of test areas. For *general accuracy*, confusion between any two spectrally neighbouring classes is allowed and considered an error that can be tolerated. For example, it is not very critical whether a pixel should belong to the class of stands with 60-70% hardwood or that with 70-80% hardwood, especially since the reference "ground truth" information also has its margin of error. This concept of general accuracy is introduced because it is felt to be more representative of typical Landsat unsupervised forestry classifications, in which a precise class breakdown is not usually attempted.

In general, the summer classes were about 45 percent accurate (strict accuracy) with up to 15 percent confusion between any two spectrally neighbouring classes (for a general accuracy of 75 percent) and 25 percent confusion with other classes. The winter classifications were 47 percent accurate on average. The strict and general accuracies obtained from the classifications are shown in Table 5. These figures were calculated using representative test samples based on the forest cover map and covered a total of 6.1% of the Lake Traverse image.

Table 5. Classification accuracies

Classification	Strict accuracy	General accuracy
SUMMER 74	44%	75%
SUMMER 76	43%	82%
SUMMER 81	43%	75%
WINTER 76	49%	
WINTER 81	44%	

The classifier also produces a *reliability image* with each classification image. The values for each pixel in the former image correspond to how reliably that pixel was assigned to a class in the latter. Commensurate with the nature of this classifier (basically a parallelepiped classifier), the reliability of each pixel is defined as a measure of the distance, in the multispectral space, between the pixel and the mode of the class it belongs to. In other words, it reflects how close a pixel is to the most common pixels of that class. This measure of reliability is independent of other classes. The unreliable pixels are those close to the boundaries of their class, independently of how close or far they might be from the boundaries or modes of other classes (Alvo and Goldberg 1981, Goldberg et al. 1982a). The classifications and their reliability images are stored in an image database and later fed one by one in chronological order to the expert system.

Testing the FEP on the Lake Traverse Area

a) The testing environment

Before running the FEP, a set of thresholds for the class and reliability transition threshold tables must be provided. This was done with the help of classification comparisons similar to the one produced by the FEP (See Table 1, where two summer

classifications are compared), but more exhaustive. That is, by looking at such classification comparisons, the user can sense the amount of changes typically occurring from one class to another and can set thresholds accordingly. Ideally, these thresholds should ensure that most artificial changes are recognized as such and considered to be non-significant transitions by the software. Unfortunately, as detailed below, this does not always happen.

For the decision-making section of the FEP to be activated, the detection section has to encounter a class difference at the pixel level between two classifications and that change of class must correspond to a transition that has been judged significant by the filtering section (see Figure 1.b). Because a filtering mechanism based on simple thresholds is far from perfect, the decision tables are actually used by the system in three possible situations:

1) *When a real change is present and the change is detected due to the quantity of the pixels affected.* A class difference is considered a real change when the number of pixels that has changed from one class to another exceeds the user-entered threshold for that transition. For drastic real changes (e.g., softwood to open) the threshold would most probably be zero. For slow real changes (e.g., open to 60% softwood, that is, regrowth), the percentage of changing pixels would have to be higher than the expected percentage of artificial differences.
2) *When enough artificial differences are present* such that these differences are considered real changes because the number of changed pixels is above the threshold for that transition.
3) *When a real change is present and the change is detected due to the high reliability of the pixels affected.* Globally, for that transition, the number of changed pixels is below the threshold in the class transition threshold table. However, more than X% of the changed pixels are more reliable in the new classification than in the old (X being the threshold in the reliability transition threshold table).

Ideally, only situations 1 and 3, implying real changes, should have reached the decision stage. In these cases, there is a legitimate need for a decision to be taken. However, due to the weaknesses of the filtering process, situations of type 2 often occur. On the surface, this should not to be a major problem because the user will presumably realize the artificial nature of the detected changes. It only forces the user to make many more decisions than would have been required normally. Nevertheless, these decisions sometimes appear irrelevant and can be hard to make. In addition, the fact that these cases reach the decision section implies that they bypass the reliability increase or decrease functions. In practice, over time, this situation is probably responsible for slightly corrupting the information (see section 6).

b) The case study - step by step

As the FEP is run, new classifications are introduced one at a time, in chronological order. The introduction of the first classification, Summer 74, is straightforward. It is compared with a null classification (the previous summer - which is unknown) and thus triggers the decision process for all pixels. Because both the old classification and the old decision image are nonexistent, it suffices for the user to create rules that will link the summer classes to corresponding decision classes in order for the system to create a new decision image. Here, the first decision image essentially duplicates the Summer 74 classification, because the decision image's forest type classes are, for the sake of simplicity, generally the same and in the same order as the classification's forest type classes (see Table 3). Similarly, the classification's reliability image becomes the first decision's reliability image.

In Decision Image 1, recent cuts are sizeable (Figure 3, in red). Note that in this case they occur because they were classified as such in the Summer 74 classification. Generally, however, if cut-over areas are not a sizeable part of an image, an automatic classification may miss them completely. In the subsequent steps of this FEP run, the identification of cut-over areas generally comes from making inferences (using the forestry rules) and is thus more significant to this case study. Notice also how the roads are confused with recent cuts because of a similarity in spectral signatures. In addition, some cut areas where a substantial number of hardwoods were left standing are noticeable by their mixture of green and red pixels.

The next step is the integration of the Winter 76 classification, producing Decision Image 2

(Figure 4) and its reliability. Similar to the previous step, this classification is compared with a null classification, which makes the system consider that real changes have occurred for all pixels (and thus, use the decision tables). However, here the net effect is that the Winter 76 classification ends up being indirectly compared to Decision Image 1 (and thus even more indirectly with Summer 74). Because of this, a lot of decision rules had to be entered at this point. Most of these rules only established the connections between the decision (or summer) classes and the winter classes, telling the FEP for example, that a transition from summer classes 6 and 7 to winter class 4 is not actually a change (all are softwood classes) and that decisions 6 and 7 should be kept. Other rules reflected true decisions that corresponded to the detection of real changes or to a gain of information over that conveyed by a single summer classification.

Below are some of the most interesting observations and decisions made. In Decision Image 1 (and Summer 74), class 5 represents the mixedwood forested areas as well as the open areas covered by grass, shrubs, and wetlands along streams, as they have very similar signatures. Addition of the winter classification to the system allowed the two areas to be differentiated, producing a purer decision 5, representing the mixedwood forest and creating decision 12, open areas (white). This was made possible because the latter were completely covered with snow in the winter, while the former were easily recognized as mixedwood forest. Also, as mentioned above, Summer 74 (and Decision Image 1) class 8 (red) contains roads and recent cut-over areas. The cut-over areas were separated and reassigned to decision 9, old cut (mauve), to illustrate the passing of time and allow new cut-over areas to be distinguishable. In fact, more cutting (region F, in Fig. 2) and the enlargement of an existing cut (region G) occurred during the winter of 1974 and 1975. They were never classified as cut-over areas in any classification because not enough pixels were involved to create a distinct (sizeable) class, but they were easily detected by the expert system (Figure 4, decision 8) as these areas went from pure softwood to open.

In general, looking at the overall effects of the integration of Winter 76, Decision Image 2 lacks detail compared to Decision Image 1. This is due to the fewer classes available with the winter classification and to the user's decisions favouring less pure forest types when, for example, something previously labelled as pure hardwood was classified as mixedwood in Winter 76 (a very common situation). This bias was based on the assumptions that winter classifications were usually good at quantifying the softwood component of forest stands and that the FEP would eventually attain the correct decision on the long run. However, considering the winter classification accuracies reported in Table 5, the user may have done better to favour the summer classification results contained in Decision Image 1. On the other hand, it is generally better to use conservative decisions, avoiding drastic effects on the system, and let the FEP establish the more accurate situation with time and additional information. This, it should be noted, does not imply that the new decision image contains degraded information relative to the old one. It only reflects the fact that the additional information (the winter classification) may in some cases only partially corroborate the old information.

Decision Image 3 (Figure 5) is produced by integrating the Summer 76 classification, which is compared with Summer 74 to detect changes. The new decisions are made based on the information from the two summers and Decision Image 2. Some of the areas cut previous to 1974 and labelled as old cut (9) in Decision Image 2 are assigned to decision 10, general regeneration (light blue), because some regrowth is detected. It is too early to know what type of regeneration is in progress. They will retain that label until the young conifers known to have been planted grow to the point where they can be identified as conifers in a Landsat classification. The newer cuts identified with the Winter 76 integration have been relabelled old cut (9), to show the passage of time.

During the fourth decision stage, the comparison between Winter 76 and Winter 81, some minor cut enlargements were also detected (again in one of the F regions). The decision image (Figure 6) also shows successful regeneration taking place in regions H and I. The detection of this regeneration is due to the presence of vegetation in these areas in the Winter 81 classification where none could be detected above the snow in the Winter 76 classification. Parts of these two regions later became labelled as softwood regeneration (dark brown) when the Summer 81 classification, which actually identified these areas as softwood, was

integrated into the system to produce Decision Image 5. By then, the predominantly jack pine regeneration was highly concentrated, 9 years old and 2 m in height (see Table 2).

A visual comparison between Decision Image 3 and Decision Image 4 shows that the mature stable parts of the forested areas of the scene are practically unchanged (also see Table 6). This implies that the added classifications are corroborating the old information. The reliability of these areas should increase with each addition of corroborating information to the system. This tends to support the hypothesis that the final decision image will be a much more accurate representation of reality with respect to non-changing areas than any single classification.

Table 6. Percentage of pixels in the same class when comparing two decision images

1st decision image	2nd decision image	Percentage
One	Two	63
Two	Three	88
Three	Four	90
Four	Five	95
One	Three	64
Two	Four	80
Three	Five	86

During the integration of the last classification (Summer 81), something interesting was perceived within the anomaly image. Remember that the anomaly image is an intermediate image produced by the FEP to show the user the location of pixels undergoing a significant transition for which there is no decision in the decision tables. A typical anomaly image has a black background speckled by occasional coloured dots, usually representing artificial changes misjudged as real (considered noise), and a few distinctive clusters showing the real changes needing the user's decisions. In this particular case, some of the noise followed patterns detectable by human observers. One pattern was associated with an old strip cut done in the 1960s. Another resulted from slow regeneration in jack pine reforested clear-cuts from the 1950s. Unfortunately, nothing could be done to use this information within the current system because the need for sophisticated human intervention at that level had not been planned.

In general, the Decision 5 Image (Figure 7) looks similar to Decision Image 4, a fact easily confirmed by direct comparison (see Table 6). The only significant differences are due to the growth of the jack pine regeneration. Some softwood regeneration in the cut-over areas was detected when certain areas changed from class 5 in Summer 76 to class 7 in Summer 81. However, some of it was not detected, as some areas only changed from class 5 to class 6. This transition is not considered significant due to the use of a high class transition threshold (to filter artificial changes). One can also extrapolate that these same areas would change from class 6 to class 7 in the next summer comparison and that this would not be detected either for the same reason. This illustrates another weakness of the present system, the difficulty of monitoring slowly changing processes like regrowth.

FEP weaknesses

The FEP was good at detecting changes in the forest due to logging. Small cut enlargements that had not been detected by single-date unsupervised classifications were easily recognized. However, the FEP was found to be only moderately good at monitoring the regrowth of conifers. Some areas with substantial conifer regrowth were detected and labelled properly, while others, with sparser and slower regrowth, were missed. This weakness has been attributed mainly to the inadequate filtering process of the FEP.

The FEP needs a filtering process because the typical inaccuracy and imprecision of Landsat classifications, especially when using very narrowly defined classes, complicate what on the surface appears to be a straightforward comparisons of new and old classifications of specific areas. Certain marginally classified pixels, found at spectral class boundaries or at forest stand spatial boundaries, have a tendency to switch from one class assignment to another when doing a time-sequential comparison of classifications. These changes, labelled artificial changes, need to be differentiated from those with real forestry-based causes because the former should not be of concern in the decision process.

In the FEP's filtering process, this differentiation is attempted with a rather simple quantity-based threshold system. Such a system is based on the assumption that major human-caused changes (e.g. logging) imply transitions between classes that are significantly different

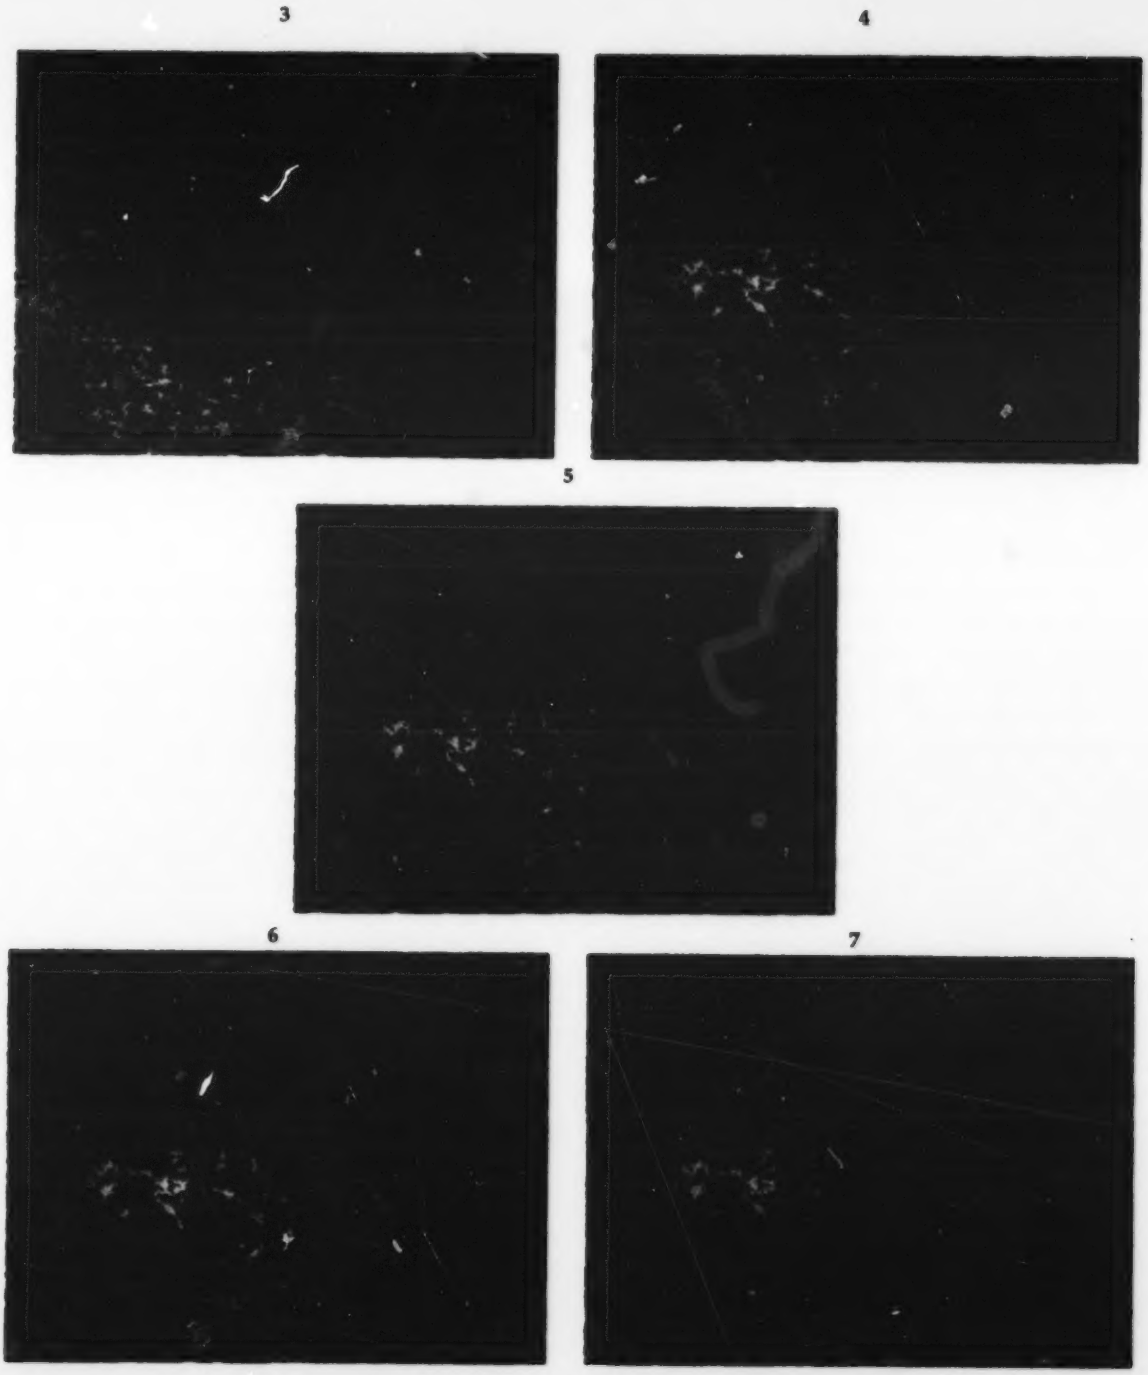

Figure 3. A Landsat image unsupervised classification of the summer of 1974 and/or Decision Image One produced by the FEP (subset of full image showing only areas where activities took place)

Figure 4. Decision Image Two which indirectly reflects the combination of the summer 1974 and winter 1976 classifications

Figure 5. Decision Image Three produced by the addition of the summer 1976 classification

Figure 6. Decision Image Four produced by the addition of the winter 1981 classification

Figure 7. Decision Image Five - produced by the addition of the summer 1981 classification

(e.g. "pure softwood" to "wide open area"), while the classes that are involved in artificial changes are more closely related (e.g. "60-70% hardwood" and "70-80% hardwood"). It is thus possible to set rather high thresholds between the latter, preventing artificial changes from reaching the decision stage, while using low thresholds for the former. When this assumption is respected, the filtering process works rather well. For example, it does not hinder the detection of logging operations. However, this cannot be said of monitoring of slow changes. For example, since regrowth generally implies a slow process and gradual transitions from class to class as the softwood becomes more apparent, its detection can be hampered by the present filtering process. The high thresholds used to ensure that artificial changes are judged non-significant changes prevent gradual changes from being detected and thus, from reaching the decision-making process.

Another weakness of the FEP is its inability to strengthen categorization (i.e. classification) of non-changing areas. Although the stability of decisions for non-changing areas increased with every additional classification (see Table 6, showing decision images becoming increasingly similar), the FEP did not necessarily settle on the proper decisions (see Table 7, showing decision images becoming increasingly inaccurate). The accuracy improvements expected with each new addition of corroborating information did not manifest themselves. Why is this?

Table 7. Decision image accuracies

Decision image	Strict accuracy	General accuracy
One	44%	75%
Two	41%	71%
Three	39%	72%
Four	36%	69%
Five	38%	72%

In the one-to-one comparison of summer classifications (Table 8) it is noteworthy that, generally, more than 50% of the pixels are not in the same class from one classification to the other. This is attributed to the use of spectrally narrow classes, which causes numerous pixels to change classes because of slight variations in class definitions (spectral shifts) or due to their spatial location (pixels at the spatial boundaries of uniform areas, whose classification may be affected by small relative misregistration of the various images). The winter classifications only differ by 33%. This is attributed to the smaller number of classes, which generally implies bigger, more uniform areas, and thus a lower number of boundary pixels. When the three summer classifications are compared, only 30.4% of the pixels never change class. When all the classifications are compared, only 12.5% of the pixels are considered similarly classified. This means that, throughout the integration of the five classifications, up to 87.5% of the pixels could make use of the decision tables at least once depending on the efficiency of the filtering process. Of course, most of these changes are artificial, because we know that at maximum only 6% of the area has undergone any change (see section 3). These artificial changes should ideally have been recognized as such by the filtering process and should never have been presented to the decision-making process.

Table 8. Percentage of pixels in the same class when comparing two classifications

1st classification	2nd classification	Percentage
SUMMER 74	SUMMER 76	46
SUMMER 74	SUMMER 81	47
SUMMER 76	SUMMER 81	53
WINTER 76	WINTER 81	67

Unfortunately, because the filtering process is based mostly on quantity thresholds and the amount of artificial changes are substantial (presumably, a situation not anticipated during the design of the FEP), most of these changes are presented to the decision-making process for evaluation. This forces the user to enter many minor and seemingly inconsequential decisions into the decision matrices. The combined effect of all these decisions appears to degrade the overall integrity of the resulting information.

An improved filtering process should not be based solely on quantity thresholds and should be able to deal gracefully with the oscillation of certain marginally classified pixels. This oscillation arises when pixels change assignment from one class to another with every new classification, typically oscillating between two spectrally neighbouring classes or at the boundaries of two spatially connected regions. It is probably sufficient for the filtering process to recognize the situation and delegate it to the reliability increase and decrease functions. Similarly, the reliability functions themselves should not be too susceptible to these minor oscillations.

From the above observations it appears that for the filtering process (and the reliability functions) to be satisfactory it must access the original spectral data from which the classifications were obtained, or at least the likelihood data in maximum likelihood classification cases. It seems that the filtering process cannot be disassociated from the classification process and dealing with classified images without their underlying spectral data may not be appropriate. In addition, whether or not the filtering process is kept at the classified information level, it must deal reasonably well with boundary pixels. With their spectral value resulting from the integration of two typical spectral reflectances, boundary pixels are hard to classify and can change class with the slightest misregistration. Consequently, the integration of a spatial operator in the filtering process may also be desirable.

CONCLUSION

A repetitive coverage makes Landsat satellite imagery well suited for monitoring and change detection, providing clouds are not a major problem. Traditionally, groups of images of an area from multiple dates were combined in various ways to make changes stand out or lend themselves to classification. A new approach consists of using single-date unsupervised classifications of images of different years and of different seasons, coupled with user identification of the classes, as input to an expert system which compares them. The present paper has described a first attempt at this approach, the Forestry Expert Package, and the experiments to test its capabilities over a forested area of Algonquin Park, Ontario.

The FEP was found to be good at detecting changes in the forest due to logging and only moderately good at monitoring the regrowth of conifers. This weakness has been attributed mainly to the inadequacy of its filtering process. A filtering process is needed to differentiate artificial changes from those with real forestry-based causes, because the former should not be of concern in the decision process. In the present FEP, this differentiation is attempted with a rather simple quantity-based threshold system. It does not hinder the detection of drastic changes, such as caused by logging operations, but it makes the monitoring of slow changes unreliable. The high thresholds used to insure that artificial changes are judged non-significant often prevent gradual changes from being detected and thus from reaching the decision-making process.

It is also the partial failure of the filtering process which has the most damaging effect on the FEP's assumed capability of rendering more accurate information on areas that have not really changed. Even though the filtering thresholds were set relatively high, some artificial changes frequently triggered the decision-making process, necessitating many minor decisions. The cumulative effect of all these decisions resulted in a final decision image with poorer overall accuracy than single-date summer classifications.

In general, a quantity-based threshold filtering system is unsatisfactory because real changes in Landsat images are usually a small proportion of the full image and can be gradual, while artificial changes are omnipresent. In future versions, the filtering process needs to be more sophisticated (for example, using spatial information).

The appropriateness of comparing only classifications of satellite imagery is also questioned. An *a priori* classification of the data imposes a somewhat artificial structure that can hinder the whole process. On the other hand, comparing classifications is a simple way of quickly eliminating most pixels (where no change occurred) from consideration in the decision-making process and it does not require that decisions be taken while considering data in multiple multidimensional spaces. These aspects should be given serious consideration before designing another expert system to monitor natural resources from successive satellite images.

Many practical and theoretical improvements are possible to make the FEP a precise and flexible management tool: a better artificial change filtering mechanism, a more flexible scheme of decision-making (variable rule syntax), access to the spectral data, a direct comparison of the new classification with the most up-to-date decision image, and more user-friendliness.

In a future version, a similar scheme could be expanded to allow the integration of thematic information derived from other sensors or obtained from existing forest inventory databases (i.e., geographic information systems). At the

Petawawa National Forestry Institute, the Digital Remote Sensing group is continuing research in these directions based on the experience acquired in this study and in related or complementary research (Leckie 1990, Kourtz 1987, 1990).

ACKNOWLEDGMENTS

The author would like to thank Peter Kourtz and Morris Goldberg for initiating this project, Don Leckie for the field information, numerous text revisions and overall support, and Alain Dombrowski, Susan Yatabe, Steve Joyce, and Mark Gillis for their pertinent comments and text revisions.

REFERENCES

Alvo, M.; Goldberg, M. 1981. A measure of reliability for classification of earth satellite data. IEEE Trans. Systems, Man and Cybernet. SMC-11:312-318.

Banner, A.V.; Lynham, T. 1981. Multitemporal analysis of Landsat data for forest cutover mapping: a trial of two procedures. Pages 233-240 *in* 7th Can. Symp. Rem. Sens., Winnipeg, Man., 8-11 September, 1981.

Gillis, M.D. 1988. Estimating change from successive static forest inventories. For. Chron. 64:352-354.

Goldberg, M.; Shlien, S. 1976. A four-dimensional histogram approach to the clustering of Landsat data. Can. J. Rem. Sens.. 2:1-11.

Goldberg, M.; Kourtz, P. 1977. The use of Landsat imagery for forest mapping: description of a proposed operational system. Pages 61-63 *in* Proc. IEEE Workshop of Picture Data Description and Management, New York.

Goldberg, M.; Alvo, M.; Karam, G. 1982a. A measure of confidence for improving classifications of Landsat imagery. Pages 155-157 *in* Proc. IEEE Pattern Recognition and Image Processing Conf., Las Vegas, Nevada, June 14-17, 1982.

Goldberg, M.; Alvo, M.; Karam, G. 1982b. Computer analysis of Landsat imagery. Technical Report by Scientific Consulting Services Inc., July 28, 1982.

Goldberg, M.; Karam, G.; Alvo, M. 1983a. A production rule-based expert system for interpreting multi-temporal Landsat data. Pages 77-82 *in* Proc. Conf. IEEE Computer Vision and Pattern Recognition, Washington, D.C., June 19-23, 1983.

Goldberg, M.; Alvo, M.; Karam, G. 1983b. The analysis of Landsat imagery using an expert system: forestry applications. Pages 493-503 *in* Proc. Sixth Int. Symp. on Automated Cartography, Ottawa, vol. II, Oct. 16-21, 1983.

Goodenough, D.G.; Goldberg, M.; Plunkett, G.; Zelek, J. 1987. An expert system for remote sensing. IEEE Trans. Geoscience Rem. Sens. GE-25:349-359.

Goodenough, D.G.; Deguise, J.-C.; Robson, M. 1990. Expert systems for DTM use in mountainous terrain. *In* ISPRS Commission VII Mid-term Symposium: Global and Environmental Monitoring Techniques and Impacts, Victoria, B.C. Sept. 17-21, 1990.

Goodenough, D.G.; Shanzer, D.; Robson, M. 1990. Knowledge-based systems for coping with clouds. *In* ISPRS Commission VII Mid-term Symposium: Global and Environmental Monitoring Techniques and Impacts. Victoria, B.C. Sept. 17-21, 1990.

Gray, S.L.; Nietmann, K. 1989. Canada's forest inventory 1986 - technical supplement. Forestry Canada, Petawawa Natl. For. Inst. Inf Rep. PI-X-86.

Hall, R.J.; Crown, P.H.; Titus, S.J. 1984. Change detection methodology for aspen defoliation with Landsat MSS digital data. Can. J. Rem. Sens. 10:135-142.

Hayes-Roth, F.; Waterman, D.A.; Lenat, D.B. 1983. Building expert systems. Addison-Wesley, Reading, MA.

Hegyi, F.; Quenet, R.V. 1983. Integration of remote sensing and computer assisted mapping technology in forestry. Can. J. Rem. Sens. 9:92-98.

Howarth, P.J.; Wickware, G.M. 1981. Procedures for change detection using Landsat digital data. Int. J. Rem. Sens. 2(3):277-291.

Kourtz, P. 1987. Expert system dispatch of forest fire control resources. AI Applications 1:1-8.

Kourtz, P. 1990. Artificial intelligence: a new tool for forest management. Can. J. For. Res. 20:428-437.

Leckie, D.G. 1990. Advances in remote sensing technologies for forest surveys and management. Can. J. For. Res. 20:464-483.

McKeown, D.M., Jr. 1984. Knowledge-based aerial photo interpretation. Photogrammetrica 39:91-123.

McKeown, D.M., Jr.; Perlant, F.P.; Shufelt, J. 1990. Information fusion in cartographic feature extraction from aerial imagery. *In* ISPRS Commission VII Mid-term Symposium: Global and Environmental Monitoring Techniques and Impacts. Victoria, B.C. Sept. 17-21, 1990.

Mooneyhan, D.W. 1983. The potential of expert systems for remote sensing application. Pages 4.1-4.5 *in* Proc. Int. Geoscience and Remote Sensing Symp. (IGARSS), San Francisco, 2(FP-1).

Moore, T.G.E.; Lockwood, C.G. 1990. The HSG wood supply model: description and user's manual. For. Can., Petawawa Natl. For. Inst. Inf. Rep. PI-X-98.

Nagao, M.; Matsuyama, T. 1980. A structural analysis of complex aerial photographs. Plenum Press, New York.

NASA. 1979. Forest resource information system. Lyndon B. Johnson Space Center, Houston, Texas.

Newnham, R.M. 1987. FIRFOR - A simple forest management model. For. Can., Petawawa Natl. For. Inst. Inf. Rep. PI-X-72.

Pilon, P.G.; Wiart, R.J. 1990. Operational forest applications using Landsat TM data: The British Columbia experience. Geocarto Int. 1:25-30.

Power, M. 1987. FIDSINFOBASE - The forest insect and disease survey information system. Can. For. Serv. Petawawa Natl. For. Inst. Inf. Rep. PI-X-65.

Tinney, L.R.; Sailer, C.; Estes, J.E. 1983. Applications of artificial intelligence to remote sensing. Pages 225-269 *in* 17th Int. Symp. on Remote Sensing of Environment, Ann Arbor, Michigan, May 9-13, 1983.